No. 1 ESS ADF:

Magnetic Tape Subsystem

By J. L. POTTER, MRS. F. B. STREBENDT and J. R. WILLIAMS

(Manuscript received May 18, 1970)

The No. 1 ESS ADF magnetic tape subsystem provides a high-volume, low-cost storage medium in which user messages are filed, and from which current message copies can be automatically retrieved upon user request. Message handling and statistical information concerning all message switching transactions are also filed on tape, providing a journal record for the switching center. Emphasis on maintenance and defensive programming techniques enables the tape subsystem to provide continuous and excellent service. This article describes the tape equipment and program package.

I. INTRODUCTION

Message switching systems designed for nationwide use must provide a high degree of transmission reliability and protection against loss of messages. Extensive efforts have been made in the design of the No. 1 ESS ADF to insure these system objectives. Once the system has accepted a message from its originator, the ADF system bas the sole responsibility for a correct and guaranteed delivery. Nevertheless, messages received at a user's terminal may be lost by an attendant, may be garbled because of transmission difficulties such as fades, distortion, and line hits, or may be lost as a result of a paper jam in the receiving terminal. To counter such occurrences, the ADF system has a message retrieval feature, whereby a missing or questionable message may be retrieved upon request. Magnetic tape is the vehicle used for long term message storage and retrieval in the ADF system.

Considerations which led to the selection of magnetic tape as the message recording medium included storage volume, access speed, and frequency of access. A copy of every message handled by the ADF system is saved for days in this message file, called the permanent file. This represents an exceptionally large storage requirement be-

cause of the high message handling capacity of the ADF system.[1] The use of magnetic tape was indicated by this need. Speed of access was found to be reasonably noncritical; however, the expected volume of retrieval requests was estimated at 1 percent of messages terminated. These considerations economically precluded the use of manual retrieval techniques and led to the use of automatic on-line magnetic tape storage and retrieval.

A key to the message retrieval process is a cross-reference file maintained in the message store.[2] The storage locations of all retrievable messages for each user who has the message retrieval feature are maintained in this file.

A second ADF system information filing requirement which is fulfilled by the magnetic tape subsystem is the recording of journal records for all message switching transactions involving the ADF system. These records, placed on magnetic tapes called journal file tape reels, provide a means for evaluating overall system performance. Further, the journal files enable telephone company retrieval center attendants to assist the ADF user with special retrieval requests. To gain access to the journal file, a journal file retrieval process has been included in the No. 1 ESS ADF magnetic tape subsystem.

The magnetic tape subsystem, consisting of tape equipment and a complete tape program package, is integrally provided for each ADF installation. Two tape unit controls with dedicated buffer control[1] tape hardware sequencers provide electrical access to and control of up to 16 tape units.[1] Maintenance software and hardware is provided to promptly detect hardware anomalies or failures.[3,4] This software provides for (*i*) quick recognition of and removal of faulty units without adversely affecting operational performance, (*ii*) diagnosis and isolation of faults to basic elements, and (*iii*) routine exercise of all tape equipment to insure functional integrity.

II. TAPE EQUIPMENT

Message and journal data storage and retrieval operations are performed under software control on tape equipment illustrated in Fig. 1. The tape control hardware is divided into two independent communities, each having a tape instruction queue in buffer control call store, a tape sequencer in each buffer control, and a tape unit control. Each community can gain access to each of the tape units in the office, provided that both communities are not simultaneously competing for the same tape unit.

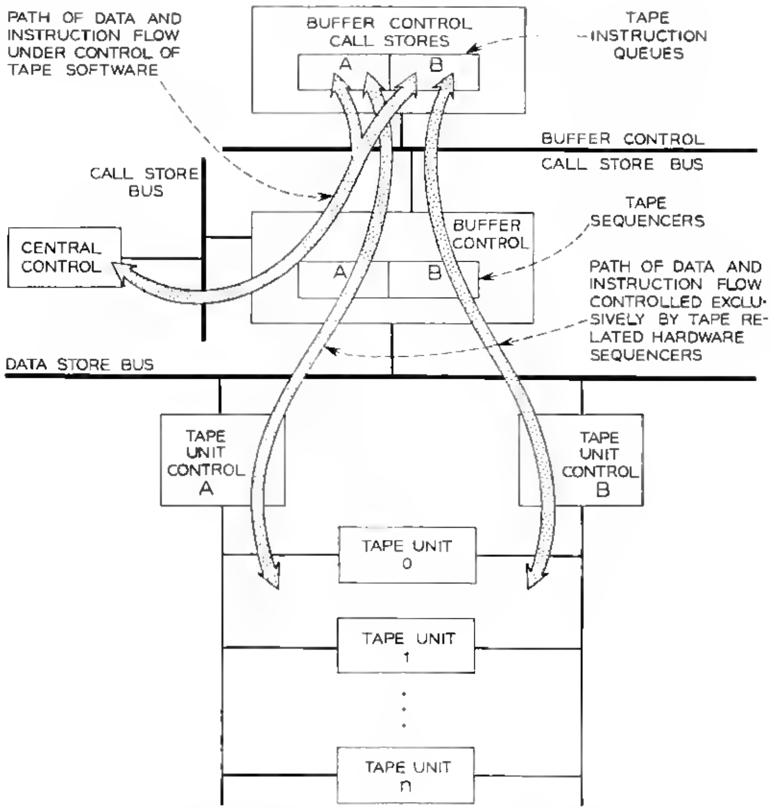

Fig. 1—Block diagram illustrating hardware units, buses, and paths of data flow involved in tape operations.

Duplication of tape related hardware units is limited to the buffer control tape sequencers and the tape queues in buffer control call store. Each buffer control has a sequencer for each tape community and thus an equipment outage will not impair tape system operation. Tape queue duplication is a consequence of buffer control call store duplication. Each tape unit control operates as a simplex unit, communicating with the buffer control tape sequencer appropriate to the configuration.

Tape instructions are loaded in the tape queues by software orders and are independently executed by the hardware. At the conclusion of the operation, completion (status) reports are placed in the tape queues for software appraisal and action.

Functionally, tape filing operations are processed on the tape unit control designated by software as primary, and retrieval operations are executed on the tape unit control designated secondary. Tape filing operations are nondeferrable, since prompt release of message store and call store facilities following final message termination is necessary to achieve the high capacity of the No. 1 ESS ADF.[5] Tbus whenever a tape unit control must be removed from service because of faults, the other tape unit control is assigned the primary function. Retrievals are deferred until both tape unit controls are in service.

Each of the tape units is automatically assigned a function appropriate to the status of the tape mounted on it. Two tape units are used for permanent message filing, the active permanent file, and the standby permanent file. Filing is automatically switched between these tape units to allow prompt retrieval of messages on the active unit. In essence, filing is accomplisbed on the active unit, while the standby unit is available for retrieval work. When the standby is inactive, its tape is positioned just after the last record written, ready to return to the active state if a retrieval is requested from the active permanent file.

When an active permanent file tape unit reaches the end of a tape, a replacement is assigned from a pool of ready tape units, that is, tape units with clean reels of tape positioned at the beginning of the tape, loaded and in service. Then the replaced active permanent file is rewound and placed in the pool of on-line tape units. These tape units provide the large bulk storage of messages which is accessed by tape software in response to message retrieval requests.

Tbe ready tape unit pool is maintained at a minimum of two tape units. If tbis pool becomes deficient, the on-line permanent file tape unit with tbe oldest reel is automatically removed from service, and an order is sent to retrieval center personnel to install a clean tape. When the new tape is loaded, the tape unit is placed in tbe ready pool. The old reel then is placed in off-line storage.

Journal records are written on a single tape unit assigned for that task. When a journal file reel bas been completely written, it is removed and placed in off-line storage. This allows no retrievals from the active journal file, but such retrievals are seldom required before the complete tape reel is written. If required, provision is made for terminating the filing assignment on the current active jounal file reel.

Retrievals from journal file and permanent file reels that are not mounted on tape units are accomplished by the assignment of an off-

line retrieval tape unit. Retrieval center personnel are directed by the program via teletypewriter to make required reel changes during the execution of these retrievals. When off-line work is completed, the off-line retrieval tape unit assignment is terminated, and the tape unit rejoins the ready pool.

Tape hardware configuration and control is automatically accomplished by tape software. Retrieval center personnel are required to change reels only as directed hy tbe system and to bandle special cases of nonretrievahle messages, reported hy the system.

To provide the features discussed, the tape hardware is capahle of executing the following operational orders:

(*i*) Read—forward direction,
(*ii*) Search—permanent file forward and reverse (hased on matching the first word in record), journal file forward (hased on matching on three-word groups within journal file records),
(*iii*) Write—forward direction of variahle length records,
(*iv*) Backspace—up to 15 records per command,
(*v*) Advance—up to 15 records per command, and
(*vi*) Rewind—at 225 inches per second.

2.1 *Tape Unit*

Each tape unit in the No. 1 ESS ADF magnetic tape subsystem includes a special design Ampex Corporation transport with associated transport servo electronics. Read amplifiers and logic for transport control and communications with the tape unit control were designed hy Bell Laboratories. Nine-channel, 800-bits per inch, nonreturn to zero operation at 56.8 inches per second is provided for normal forward and reverse operation. Rewinding at 225 inches per second enables a 2400-foot reel to he rewound in two minutes. Figure 2 shows a typical tape unit installation.

The tape unit responds to several commands from a tape unit control in order to execute the various instructions: write, read, forward, reverse, move, stop, rewind, select and unselect. Data is written on tape in records of 1 to 15 hlocks with 3/4 inch gaps hetween records. Each block consists of 31 call store words of data plus one link word. A word is 24 hits long, with 23 information hits and one parity hit. On tape, each word is written in three segments called hytes. A byte is nine bits, eight information hits and one parity bit.

To insure integrity of data written on tape, eacb tape unit is

Fig. 2—A typical tape unit installation.

equipped for read after write, thus enabling the tape unit control to verify each write operation.

2.2 *Tape Unit Control*

The tape unit control is a consolidation of common tape control equipment that interfaces the buffer control and tape unit. Extensive circuitry is provided for error detection in tape operations. Figure 3 pictures the tape unit control.

The tape unit control receives tape instructions from the buffer control, stores each instruction in its instruction register, decodes the instructions, and controls tape unit operation during execution of the instruction. Data transfers between tape and call store pass through the tape unit control's 24-bit data register, where word-to-byte transformations are accomplished.

Read, write, and search control circuits coordinate data transfers between tape units and the buffer control, and control operation of the selected tape unit.

2.2.1 *Operation*

Upon receipt of an instruction, the tape unit control decodes the order and selects the proper tape unit. For a write operation, the tape unit control orders the selected tape unit into the write, forward, and

Fig. 3—Tape unit control.

move states. Data is transmitted to the tape unit, a byte at a time, at a rate necessary to achieve 800-bit-per-inch density at 56.8 inches per second. When the order is completed, the tape unit control stops and releases the tape unit, and reports status to the buffer control. Read instructions are similarly handled by the tape unit control.

At the beginning of the search operation, search data received from the buffer control is saved in the tape unit control. Next, the tape unit is selected and controlled as in a read operation. Data read from tape is not sent to the buffer control; instead, the search control circuitry matches search data within each record obtained from tape against the search data received from the buffer control. When an exact match is obtained, the search operation is terminated, the tape unit is stopped and released, and successful search status is reported to the buffer control.

Two types of search instructions are provided. The journal file search instruction searches the first five blocks of every record to match 9-byte search entries. Section 4.4 of this article describes the layout of the journal file record. The permanent file search matches a 24-bit search number against the first word (three bytes) of every record passed on the tape. Permanent files can be searched in either forward or reverse.

2.2.2 *Integrity*

The tape unit control contains several mechanisms for guaranteeing the integrity of data transmission and storage. One method used is the all-seems-well response[6] to buffer control transmissions. Absence of an all-seems-well response causes a system interrupt, which results in immediate program action towards solving the difficulty.

Communications between the tape unit and its control are monitored via status registers in each tape unit control. Failures or data errors cause unique status bit combinations in the status register. At the conclusion of every operation the contents of the status register are transferred to the tape queue in buffer control call store. Program detection of a bad status triggers the tape system fault recognition software.

Transverse and longitudinal parity checks insure the integrity of data stored on tape. As a part of the write operation, the tape unit control generates parity over each 8-bit byte and stores it as the ninth bit of the byte. During a read operation, after parity has been checked, the tape unit control strips off this parity bit. A parity

failure causes the tape character error counter in the status register to be incremented and the tape character error indicator to be set.

As data is being written, the tape unit control also computes longitudinal parity over the entire record. At the end of the record, this byte is written on tape. This serves as an additional error detection device used in read operations.

The tape unit sends data to the tape unit control in an amplified analog form over a common bus. For reliability, the tape unit control has two signal detectors, the high and the low amplitude threshold read registers, listening to this bus. If both registers agree, the information is accepted as valid. However, if the low register's parity passes and the high register's parity does not, the data is judged in error and the tape character error indicator in the tape unit control is set. Nevertheless, the entire record is read. This enables the No. 1 ESS ADF to retrieve the best copy of the requested data, if after multiple attempts parity errors continue to occur.

The status report, initiated by the tape unit control upon completion of an operation, enables the software to determine the success or failure of the order. Decisions to try again, remove equipment from service, or continue normal operational sequence are heavily based on the status report. The layout of the tape status report (Table I) provides insight into the error detecting and reporting scheme of the tape unit control.

2.3 *Buffer Control Tape Sequencers*

Within each buffer control, dedicated tape sequencers are provided for each of the two tape unit controls. These sequencers control the flow of information between the corresponding tape unit control and its dedicated tape instruction queue in buffer control call store.

Each tape instruction queue consists of 32 buffer control call store words. Tape programs place the instructions, in two-word instruction sets, in the desired queue. The first word of the instruction contains the address of the buffer control call store link list of data blocks to be used for the data transfer. The second word contains the operation code (read, write, backspace, and so on), the tape unit involved, and other pertinent data.

The sequencer is responsible for obtaining the data from (or storing data in) buffer control call store and sending it to (or receiving it from) the tape unit control upon request. Data in buffer control call store is stored in 32-word block link lists. The sequencer inter-

Table I—Layout of Tape Status Word

Bit	Description
0	Beginning of tape
1	End of tape
2	Rewinding
3	Unsuccessful search
4	Tape unit failure in setup
5	Tape unit failure in run
6	No character response
7	Data time-out
8	Parity error
9	Word count error
10	Parity generator error
11	Longitudinal parity mismatch
12	Write longitudinal parity mismatch
13	Longitudinal parity error
14	Tape character error
15	Counter overflow
16 17 18 19	Word error counter
20	Buffer control invalid sequence
21	Successful operation
22	Successful operation

prets the last word of each block as the link address of the next block of data. Each transfer of data from the buffer control to the tape unit control is verified by all-seems-well techniques.

The tape sequencer responds to instruction, data, and status requests from the tape unit control. In the idle state, a tape unit control continuously sends instruction requests until the sequencer finds work to do. When an instruction is found in the tape instruction queue, the sequencer sends the instruction to the tape unit control. If a transfer of data to tape is involved, the tape unit control sends a data request to the tape sequencer, in response to which the sequencer obtains a data word from buffer control call store and transfers it to the tape unit control. When this data has been transferred to tape, the tape unit control sends another data request to the sequencer. The next word is sent to the tape unit control, and the process continues until all data has been transferred to tape.

When the tape unit control has completed execution of the instruction, it sends a status request to the buffer control tape sequencer.

In reply to this, the tape sequencer fetches the execution status from the tape unit control and stores it in the queue for later software examination. Unsuccessful status freezes the sequencer, thus retaining tape unit control and sequencer information for fault recognition processing. After a successful operation, the sequencer proceeds to the next instruction in the tape queue.

Checks are made within the tape sequencer as well as the tape unit control to insure that the correct number of data words have been transmitted. If errors are detected, appropriate status bits are set and the sequencer is stopped as for other errors.

III. MAINTENANCE CONCEPT

The objective of the tape subsystem maintenance scheme is to provide a high degree of dependability and maintainability, consistent with ESS standards. To meet this objective, sophisticated techniques are used in fault detection and recognition, diagnosis, error analysis, and routine exercise.

3.1 *Environment*

Although duplicated hardware matching is a standard technique used in high reliability systems, it is not used in all segments of the tape subsystem. Notably, the tape units and tape unit controls used in the No. 1 ESS ADF magnetic tape subsystem are operated in a simplex mode. Reliable simplex operation of these units is provided by extensive maintenance hardware and software for detecting, reporting, and analyzing errors. The simple fact that two tape unit controls and several tape units are installed in the No. 1 ESS ADF switching center does mean that functional redundancy effectively exists. All critical filing operations are continued, for example, when a tape unit control is out of service. For such an occurrence, the retrieval function is simply deferred until the second tape unit control is restored. The powerful ESS diagnostic approach used for tape subsystem equipment insures that equipment downtime is minimized.

Tape maintenance procedures must deal with several types of faults and abnormalities. The hard logic failure is easily detected and isolated, but intermittent logic and mechanical problems are a serious challenge. Numerous retries of failing instructions and error analysis methods are used to overcome the latter problems.

3.2 *Fault Recognition*

The tape subsystem fault recognition software:

(*i*) Determines action required when tape subsystem errors are detected.

(*ii*) Identifies true faults from superficial failures such as minor tape imperfections.

(*iii*) Isolates true faults to the defective equipment unit.

(*iv*) Removes faulty equipment from service for diagnosis and repair.

Fault recognition processing begins upon detection of tape hardware interrupts or failing status reports.[3,4] Faults which cause interrupts are isolated by the program to tbe faulty tape unit control, buffer control, or data store bus. Tbe faulty unit is then removed from service and diagnosed.

When fault recognition is stimulated by an unsuccessful status report in the tape instruction queue, the failing instruction is repeated up to four times. If any of these four attempts passes, no further action is taken. After four consecutive failures the instruction is tried on the other tape unit control. A success implies that the original tape unit control was at fault. If no attempt succeeds, the tape unit is assumed to be at fault since it cannot work with either tape unit control.

This simple algorithm bas some complications resulting from the uniqueness of the tape system. First, it is difficult to tell if the tape unit control, the tape unit, or tbe magnetic tape is at fault. For example, a longitudinal parity failure during a tape operation may be caused by defective tape or a bad parity circuit in the tape unit control. A simplification is made by assuming that tape imperfections are minute and isolated to a small portion of tape. For a failing write instruction, before trying the instruction on the other tape unit control, up to six incbes of tape is skipped and the write process repeated. If this fails, tbe failure is assumed to be an electronic fault. Success implies a tape imperfection, and no further action is taken.

When a tape unit is removed from service, the tape is placed in off-line storage. However, wben a tape unit control is removed from service, the tape unit it was addressing is not removed. Incorrect or incomplete records are erased from tbe tape using the in-service tape unit control. Correct data is then written on the tape using the in-service tape unit control.

3.3 *Error Analysis*

Error analysis provides the means by which tape system fault recognition copes with marginal electronic circuitry and mechanical ab-

normalities. Both prohlems are usually intermittent and not consistently reproducible. Therefore records are kept in call store of all failures. When any one type of failure is excessive the unit is removed from service and diagnosed.

3.4 *Diagnostic*

The tape suhsystem diagnostic software provides a thorough test of all tape-related hardware equipment. The objective, of course, is to isolate failures to individual replaceable circuit elements such as circuit packs. Diagnostic procedures are provided for the tape unit controls, tape units, tape sequencers, and connecting buses.

The diagnostic software reports results via the maintenance teletypewriter to maintenance personnel. If an error is detected, a unique 12-digit trouble number is reported which, when checked in a trouble locating manual, identifies the faulty circuit.

Mechanical faults are difficult to diagnose. Inconsistent diagnostic results often occur, yet optional raw data diagnostic printouts assist maintenance personnel in locating faulty components.

3.5 *Exercise*

The tape subsystem exercise routinely checks idle hardware to insure that operationally undetectable faults do not remain undetected. This is accomplished by periodically running all equipment diagnostics on those units which have not been diagnosed within a given time. This technique guards against faults in the maintenance failure detecting circuitry which could mask other faults, and thus prevent error detection.

Each tape unit control is routinely exercised, provided that one tape unit control remains in service at all times. All ready tape units are exercised. However, exercises of tape units assigned to an active filing, retrieving, or on-line storage pool are deferred until termination of the functional assignment.

If an exercise passes, the equipment unit is returned to service. If a fault is detected, a diagnostic report is sent to the maintenance personnel, and the equipment unit is left out of service.

IV. OPERATIONAL FEATURES

The primary purpose of the magnetic tape subsystem is message retrieval. The permanent filing feature exists mainly to support message retrieval. As Section II discusses, on-line permanent file tape

units provide for the bulk storage of customer messages. Assuming a complement of 16 tape units, 10 would normally be in the on-line pool. Using 2400-foot reels and average message statistics, the on-line retrieval capacity is approximately 125,000 messages. Older messages are retrievable semi-automatically from off-line reels.

The configuration, control, and use of the tape subsystem equipment is totally automatic. The system maintains in the cross-reference file the storage locations of all retrievable messages handled by the system. Further, the location of every tape reel is maintained in a call store table. Both of these are routinely audited to insure integrity. Retrieval center personnel need change tape reels only as indicated by the system. For permanent and journal filing as well as on-line retrieval operations, no manual intervention or tape handling is required.

Retrieval center personnel have access to the journal records via the journal file retrieval process. Information filed on the journal tape gives a history of message switching transactions but does not include any message text. Through off-line analysis of the journal tapes, telephone company personnel can ascertain the quality of total user service, operating and traffic statistics, and switching transactions. The journal file retrieval capability also enables retrieval center personnel to assist the user with special problems for any given message.

4.1 *Permanent Filing*

Each message handled by the No. 1 ESS ADF is recorded on permanent file tape after all terminations for the message have been completed.

For program and tape packing economies, messages are divided into three categories, based on total file length (that is, the total number of message heading, text, and message processing blocks.) Messages of five blocks or less (short messages) are packed two or three to a record, provided that sufficient messages are terminated with this characteristic within a fixed time. Since the maximum tape record is 15 blocks, messages greater than 15 blocks (long messages) are filed in several records. Messages from 6 to 15 blocks long (normal messages) are not packed and thus record size in blocks is equivalent to message length.

In the first word of each tape record is placed a unique 23-bit search number. As Fig. 4 shows, this number consists of a base search number, and long and short message indexes. The short message index is set to zero in the tape record. After a message has been filed on

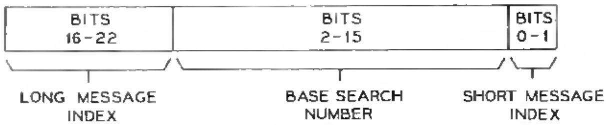

Fig. 4—Breakdown of search number used to locate a message, or part of a message, on tape.

tape, the base search number (and short message index if applicable) and tape reel number are placed in the cross-reference file forming the basis for searching if a retrieval is requested. When permanent filing is started on a new tape, the base search number is set to one. For subsequent records that are not long message residue segments, the base search number is increased by one, providing the uniqueness.

For a normal length message, the long and short message indexes are always zero. For a long message, the long message index is zero for the first record and is increased by one for each subsequent record of the same message. The base search number remains the same for all records of a given long message. For short messages, one base search number applies to all messages within a record. The short message index which is placed in the cross-reference file entry identifies where the message is packed within the record. This index is zero in the search slot at the beginning of the tape record.

Permanent file strategy does not require that all records of the long message be written consecutively. High efficiency is obtained by involving the program is concurrent processing of numerous messages. Figure 5 shows an example of final tape data packing.

A tape-to-tape link is provided when the end of tape is reached when partially recorded long messages exist. The balance of these messages are written on the replacement permanent file tape, and a link record is written on the first tape, linking each such message to its new tape address. Thus a tape can be completely used without sacrificing retrieval capabilities. At the same time, call store and primary message store space can be released when each long message segment is written on tape.

The permanent file message tape copy contains a complete copy of the message, its heading, and associated message processing blocks.[7] Thus when a message is retrieved from tape, the retrieval program readily verifies correctness and completeness, reconstructs the message, and nominates it for delivery.

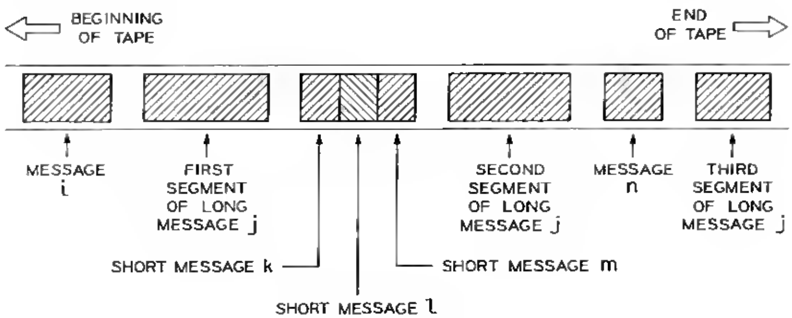

Fig. 5—An example of data packing on permanent file tape.

4.2 *Cross-Reference File*

A vital link in the message retrieval process is the cross-reference file, which is a station-by-station record in the message store of messages sent or received. The cross-reference file is the basis for message retrieval since it contains the location or retrievability status of messages handled by the No. 1 ESS ADF. The fixed cross-reference file area for a particular station can be found in that station's translation information. From this, the entry concerning a given message can be obtained via an indexing scheme using the message number.

The cross-reference file organization consists of four different kinds of blocks, each 32 words long. (See Fig. 6.) Recycling blocks point to summary blocks, and summary blocks point to file blocks. Accumulating blocks accumulate cross-reference data, and when full, their contents are transferred into file blocks.

Within each station's fixed area of cross-reference file are two recycling blocks, one for terminated messages and one for originated messages. A recycling block is an overall record of received or sent messages. The first 21 words of the recycling block are divided into seven sections of three words each. Each section represents a group of 1000 message numbers called a cycle. Message numbers for a station generally proceed from 1 to 999 and then go to 1 again, although stations may reset their message number counter before normally recycling. Tbe newest cycle is at the top of the recycling block and the oldest at the bottom. Since the seventh cycle is needed for updating, each station's cross-reference file can contain information for a maximum of 6000 messages each for send and receive message numbers.

For each cycle there is a summary block address, the date and time

of the first message number in that cycle, and the message number upon which the cycle ended. The last two words of the recycling block contain the oldest valid message number cycle and the oldest message number therein for which cross-reference file information has not yet been released.

Summary blocks and file blocks are seized from a pool common to all stations. The number of these blocks that a station may seize is a subscriber feature directly affecting retrieval capability. Thus, a station may elect to be able to retrieve from none to the latest 6000 messages. A summary block contains the addresses of all the file blocks pertaining to a given cycle within a recycling block. Each of the 32

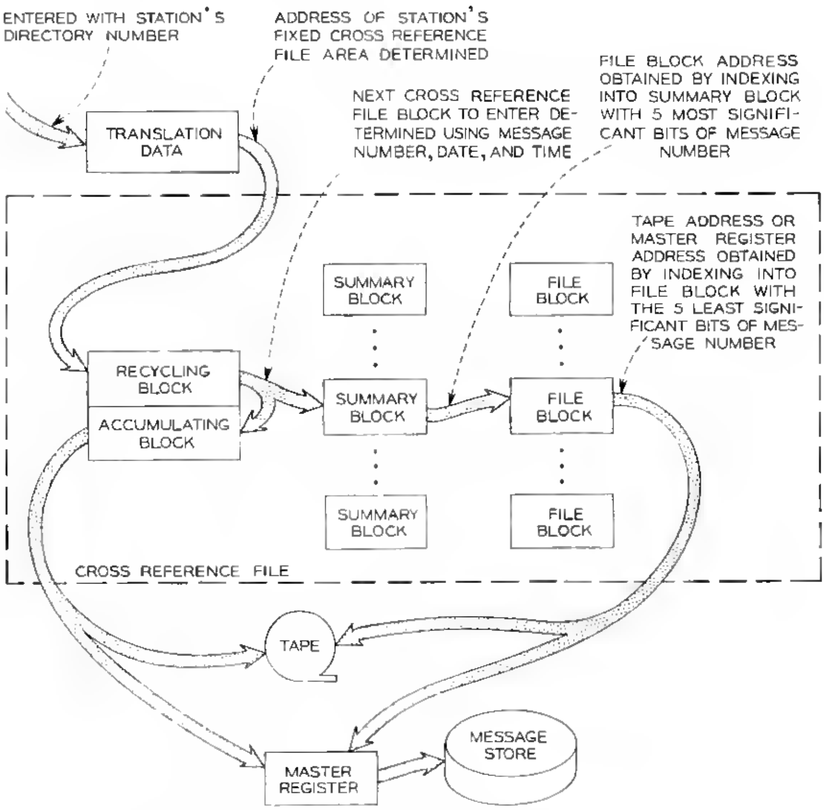

Fig. 6—Layout of cross reference file, illustrating indexing scheme for locating a message on tape or in primary message store.

words of a summary block points to a file block if this is a full cycle of 1000 messages. A file block contains the location or status of a group of 32 messages such as those for message numbers 1 to 32, 33 to 64, and so on.

An accumulating block collects entries for the most current group of 32 messages. When a message number which is a multiple of 32 is reached, the contents of the entire accumulating block are transferred to a file block. The appropriate slot within the summary block is also updated with this file block address. Accumulating blocks, like recycling blocks, are part of each station's fixed cross-reference file area. There is one accumulating block for receive message numbers and one accumulating block for send message numbers.

As soon as a message enters the No. 1 ESS ADF, an initial entry in the cross-reference file is made for the originating station's send message number. As each message termination is completed, an initial cross-reference file entry is made for that addressee's receive message number. In these initial entries, a master register[7] address is inserted into the appropriate message number slots in the corresponding station's accumulating blocks. At this time the master register points to the address of the message in the message store. After all deliveries are made and the message is permanently filed on magnetic tape, the appropriate cross-reference file slots for the originator and all terminators are updated with the tape address of the message. If the message was not completely filed, if it was an action request, or if it was aborted in origination or delivery, the cross-reference file entries are marked with a non-retrievable code.

4.3 *Message Retrieval*

The message retrieval service feature permits user stations to retrieve their messages in the event of garbled copy, message numbers out of sequence, a station out of paper, or the like. The location of each message requested is obtained via the cross-reference file. Each message is found, processed, and sent to the retrieval destination. Examples of a garbled message, a message retrieval request for that message, and the message retrieval output are shown in Fig. 7.

4.3.1 *Requesting Message Retrieval*

The input stimulus for message retrieval is an input message called a resend action request, whereby a user station requests the No. 1 ESS ADF to resend one or more messages. Basically, the only items

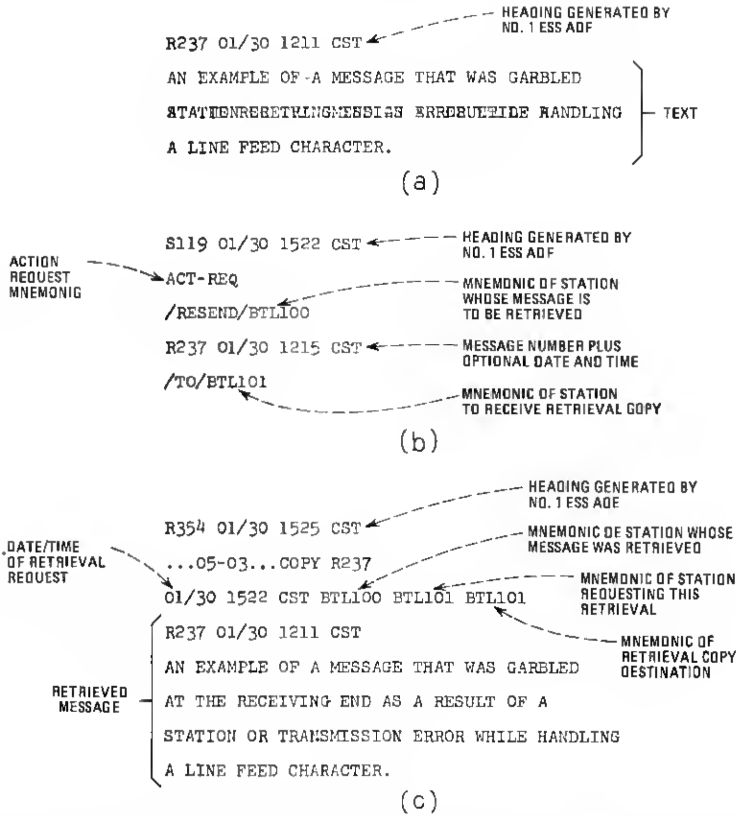

Fig. 7—(a) A garbled message received at station BTL100. (b) Message retrieval request originated at station BTL101. (c) Retrieved message copy received at station BTL101.

that must be specified are the identity of the station whose messages are to be retrieved and the message numbers desired. (A maximum of 100 message may be requested in a single action request.) Unless dates and times of the messages requested are also included, it is assumed that these messages are in the most recent cycle of 1000 messages. A variation in format permits retrieval from a given message number through the most recent message number for that station. The telephone company retrieval center has the further capability of requesting a message whose tape address is already known. In all cases, the retrieval destination, unless otherwise specified, is the action request originator.

4.3.2 *Locating Messages Requested*

When a message retrieval action request is received, the cross-reference file of the station whose message is to be retrieved is interrogated for each message number desired. First it is determined whether the message is retrievable; if it is, its location is obtained.

It could happen that a desired message is too old, that is, its pertinent cross-reference file area has been released. In this event, a service message is sent to the retrieval center, where the attendant may, through journal file retrieval, obtain the tape address, which then is used for the retrieval request.

A message number may be nonexistent or nonretrievable. An example is a requested message number larger than the most recently used message number as indicated by the message number counter. This results in a service message to the action request originator.

If a retrieval is requested for a message which is being filed, a delay mechanism causes the cross-reference file to be interrogated later.

4.3.3 *Queuing Tape Search Requests*

In most cases the cross-reference file entry for a message number contains a permanent file tape address, consisting of a tape reel number and a search number. Because of the time-consuming nature of tape operations and the fact that only one search can be performed at a time, this tape address is used to make an entry on a tape search queue. A tape search queue is a call-store-linked list of retrieval requests for a given tape reel. Each tape search queue entry consists of a four-word call store facility called a retrieval register. Each retrieval register contains the message number requested, the tape address of the desired message, and a reference to the action request (which is stored in primary message store until the message has been found).

Several tape search queues are involved in message retrieval. For on-line tape reels there is a queue per tape unit. For off-line tape reels there is only one queue. Entries on any tape search queue are ordered according to ascending search number. In the case of the off-line tape search queue, entries are grouped first-in first-out, according to tape reel number. Within each off-line tape reel number group, entries are further arranged in order of search number. Tape searching normally proceeds in the direction of increasing search numbers.

4.3.4 *Scheduling Tape Searches*

All tape search queue entries are serviced in reasonable time because

of a tape search priority scheduling scheme. New tape search operations are started on the highest priority scheduled tape search queue which has entries.

The four categories of tape search queues, listed in order of priority, are:

(*i*) *Ready-On-Way:* For a tape unit with the oldest completed on-line permanent file tape scheduled for reel change on completion of search.

(*ii*) *Timed Out:* For an on-line tape search queue that has not been serviced for an excessive time.

(*iii*) *Off-Line:* For search requests for either permanent or journal file off-line tapes. Off-line permanent file searches are scheduled to alternate with journal file retrieval searches.

(*iv*) *On Timing:* When an on-line tape search queue is initially set up, it is given this lowest priority. A timer is also set so that the priority can be upgraded if this tape search queue is not processed within a reasonable time. If a search is to be performed from this schedule, the longest queue is selected first.

4.3.5 *Searching A Permanent File Tape*

The prerequisites for starting a search are that (*i*) another search is not in progress, (*ii*) the secondary tape unit control is available, and (*iii*) the tape unit (corresponding to the tape search queue scheduled for processing) is available for searching.

If the tape selected for searching is on the active permanent file tape unit, writing of the permanent file is first switched to the standby permanent file tape unit. The former active permanent file tape unit becomes the standby, and its tape is rewound and made available for searching. When retrieval is completed on the standby, its tape is positioned just after the last record written so that it may immediately resume permanent file writing when needed.

If an off-line permanent file tape is scheduled for searching, an off-line tape unit is assigned. The tape attendant is instructed to replace its tape reel with the desired off-line permanent file tape. When this is accomplished, the tape unit is made available for searching. When retrieval is completed from this off-line tape reel, the tape attendant is requested to put the next tape on.

Searching a particular permanent file tape begins with processing the first entry on the corresponding tape search queue. A forward search is started for the search number contained in the first entry. When the proper record is found, and the tape backspaced for read-

ing, buffer control call store tape blocks are seized for receiving the data, and a read is ordered. When the read is accomplished, the message is verified for station directory number and message number. A new message is built, using the message found and an indication that this is a retrieval copy. This message is then sent to the retrieval destination.

If the message requested is long (stored in more than one record) the search-backspace-read sequence is repeated for each segment until the entire message is found. If it is a short message packed in a record with one or two other short messages, the desired message is extracted from the record and processed.

These tape operations and the subsequent processing of the message found are repeated for each entry on the tape search queue of the tape being searched. Since retrieval queue entries are grouped according to ascending search number, all entries on a tape search queue can be processed in one efficient pass along the tape. At the conclusion of the search, the tape unit is returned to the proper state and search scheduling logic is entered to select the next tape search queue for processing.

4.3.6 *Retrieving From Message Store*

A master register address in the cross-reference file slot for a requested message number implies that the message is recent. It has not yet been filed on magnetic tape and is still located in the message store. The message is immediately retrieved from message store and sent to the retrieval destination.

4.4 *Journal File*

Selected data from each message is filed in journal file records on tape. Originating and terminating directory numbers, delivery status, dates, times, tape address of permanent file message copy, and message numbers are examples of this selected information.

All journal tape records are 15 blocks long, and each record is divided into 50 entries. Each entry is composed of a three-word search slot and a six-word text slot as illustrated in Fig. 8. A three-word search slot with its corresponding six-word text slot is used for each message originator. Similarly, each message terminator is assigned search and text slots. Thus a message sent to three addresses requires a total of nine words for the originator and 27 words for the terminators.

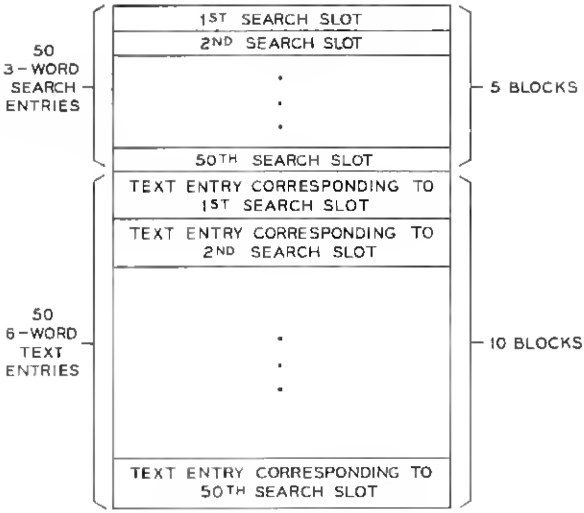

Fig. 8—Layout of data in the journal file tape record.

All information packed into the journal file record is obtained from the message processing block complex for each message. Message data is accumulated in a 15-block journal file buffer in buffer control call store. Data from individual messages is packed into the journal file buffer. When the buffer is full, its contents are written as a single record on the journal file tape.

The journal file record, unlike the permanent file record, in no way has a one-to-one relationship with the individual message. Rather, the journal tape record usually contains entries for many messages, one or two of which may overlap into adjacent records. The originator and terminators of an individual message always occupy contiguous slots in the record.

The search slot for each message originator and terminator contains the directory number, message number, and originator or terminator flag applicable to the message sender or receiver. The information contained in the search slot is unique to this message and thus forms the basis for search in a journal file retrieval. The corresponding text slot contains statistical information concerning the handling of the message by the No. 1 ESS ADF, such as delivery status, date, and time.

No. 1 ESS ADF journal records are useful because:

(*i*) They contain tape addresses of all permanent file message copies, and tbus a message may be retrieved even after the cross-reference file entry has been overwritten in the message store. A journal file retrieval process is provided for this purpose.

(*ii*) Extracts of message processing records can be retrieved from the journal file tape by telephone company retrieval center attendants for assisting users and operating center personnel with special problems.

(*iii*) Statistical studies of operating performance, traffic levels, and specialized problems are reduced to final form by analyzing the journal records on commercial data processing systems. A journal file analysis program geared to a commercial computer is operational and is frequently used for this purpose.

4.5 *Journal File Retrieval*

The journal file retrieval service feature enables retrieval center personnel to obtain origination, termination, and processing information concerning messages handled by the No. 1 ESS ADF.

4.5.1 *Requesting Journal File Retrieval*

A journal file search is initiated by an input message called a journal file resend action request. This action request specifies the directory number of the station that sent or received the message, the message number, the journal file tape reel number, and the date and time that the recording was started on the journal file tape. The journal file tape reel to be searched must be off-line. Each journal file search action request implies a searcb over the entire reel for all entries pertaining to the message specified.

4.5.2 *Queuing Journal File Search Requests*

Since journal file retrieval requests cannot be processed immediately and since each search consumes considerable time, a journal file search queue entry is made for each searcb request. As in message retrieval, a four-word call store facility called a retrieval register is used. A new entry is added to tbe end of the queue, the queue consisting of a one-way link list of retrieval registers.

4.5.3 *Scheduling Journal File Tape Searches*

Journal file search scbeduling is performed with the message retrieval tape search scheduling; journal file searches are alternated

with permanent file off-line searches. The same conditions and procedures are followed as in setting up a permanent file off-line search.

4.5.4 *Searching A Journal File Tape*

When a journal file search is next in the search priority, a journal file forward search is set up. Unlike a message retrieval search on one word (the search number), a journal file search operation is based on three words. The hardware matches on three-word entries within the search section (tbe first five blocks) of each 15-block record on the journal file tape. The first two words include tbe originator or terminator flag and a ten-digit directory number. The third word is the message number.

Processing the journal file record that is found involves locating the search slots for all the entries associated with the entry for which the search was made. The originator entry and all terminator entries for this message are found. Since the entries for a given message may extend beyond one record, additional tape operations may be required. The contents of the text slot corresponding to each search slot is then put into output format, one entry at a time, starting with the originator entry. Using this newly formed text, a journal file retrieval output message is built and sent to the retrieval center.

Such a message contains data concerning origination, termination, and handling of the message. Originator information includes the originator's directory number, the message number, the date and time of the origination, the originating message customer identity, and tbe input message status. The termination data for each terminator consists of the directory number, the message number, tbe terminating date and time, the delivery status, and the delivery precedence. The tape address of the complete message copy and the total message length are included as handling information.

After a match has been made and the record processed for output, another search is set up for the given directory number and message number. The search-read-and-process sequence is repeated until all entries have been found on the tape for tbe given directory number and message number combination. Wben tbe end of tbe tape is reached, the tape is rewound and reel change ordered, as appropriate for remaining off-line retrieval work.

V. ACKNOWLEDGMENTS

The authors wish to acknowledge the dedicated efforts and tech-

nical expertise of the many individuals at Bell Laboratories who have contributed to the development of the magnetic tape subsystem. The mutual cooperation and achievements of the numerous departments involved in the design of circuits and equipment, maintenance methods and programs, and operational software enabled the successful development of this system. Particularly noteworthy are the design achievements of Messrs. W. G. Reichert, Jr. and G. A. Van Dine for the tape unit control and tape unit hardware circuits, respectively.

REFERENCES

1. Ewin, J. C., and Giloth, P. K., "No. 1 ESS ADF: System Organization and Objectives," B.S.T.J., this issue, pp. 2733–2752.
2. Aitcheson, E. J., Ault, C. F., and Spencer, R. G., "No. 1 ESS ADF: Message Store—A Disk Memory System," B.S.T.J., this issue, pp. 2887–2913.
3. Aitcheson, E. J., and Cook, R. F., "No. 1 ESS ADF: Maintenance Plan," B.S.T.J., this issue, pp. 2831–2856.
4. Downing, R. W., Nowak, J. S., and Tuomenoksa, L. S., "No. 1 ESS Maintenance Plan," B.S.T.J., *43*, No. 5, part 1 (September 1964), pp. 1961–2019.
5. Genke, R. M., Harding, P. A., and Staehler, R. E., "No. 1 ESS Call Store—A 0.2-Megabit Ferrite Sheet Memory," B.S.T.J., *43*, No. 5, Part 1 (September 1964), pp. 2147–2191.
6. Connell, J. B., Hussey, L. W., and Ketchledge, R. W., "No. 1 ESS Bus System," B.S.T.J., *43*, No. 5, Part 1 (September 1964), pp. 2021–2054.
7. Kienzle, H. G., Nicodemus, K. L., Smith, M. T., Jr., Weber, E. W., and Zydney, H. M., "No. 1 ESS ADF: Message Processing Program Organization," B.S.T.J., this issue, pp. 2753–2829.